ECS Configuration Change Request

1. Originator	2. Log Date:	3. CCR #:	4. Rev:	5. Tel:	6. Rm #:	7. Dept.
Robert Cole	11/7/02	02-0981	-	301-925-0799	2110C	SE/HW

8. CCR Title: Add 10x180GB drive shelf to NSIDC SAN

9. Originator Signature/Date	10. Class	11. Type:	12. Need Date:
/s/ Robert Cole 11/6/02	II	CCR	11/30/02

13. Office Manager Signature/Date	14. Category of Change:	15. Priority: (If "Emergency" fill in Block 27).
/s/ James Mather 11/6/02	Update ECS Baseline Doc.	Routine

16. Documentation/Drawings Impacted *(Review and submit checklist):*	17. Schedule Impact:	18. CI(s) Affected:
920-TDN-001, 922-TDN-044		SANHW

19. Release Affected by this Change:	20. Date due to Customer:	21. Estimated Cost:
5A		None - Under 100K

22. Source Reference: ☐ NCR (attach) ☐ Action Item ☐ Tech Ref. ☐ GSFC ☐ Other:

23. Problem: (use additional Sheets if necessary)
A 10x180 GB drive shelf has been procured for the NSIDC SAN and needs to be installed.

24. Proposed Solution: (use additional sheets if necessary)
See next sheet.

25. Alternate Solution: (use additional sheets if necessary)
none

26. Consequences if Change(s) are not approved: (use additional sheets if necessary)
SAN space cannot be increased

27. Justification for Emergency (If Block 15 is "Emergency"):

28. Site(s) Affected: ☐ EDF ☐ PVC ☐ VATC ☐ EDC ☐ GSFC ☐ LaRC ☒ NSIDC ☐ SMC ☐ AK ☐ JPL ☐ EOC ☐ IDG Test Cell ☐ Other

29. Board Comments:	30. Work Assigned To:	31. CCR Closed Date:

32. EDF/SCDV CCB Chair (Sign/Date):	**Disposition:** Approved App/Com. Disapproved Withdraw Fwd/ESDIS ERB Fwd/ECS
33. M&O CCB Chair (Sign/Date): /s/ G. G. Gavigan 11/12/02	**Disposition:** Approved App/Com. Disapproved Withdraw Fwd/ESDIS ERB Fwd/ECS
34. ECS CCB Chair (Sign/Date):	**Disposition:** Approved App/Com. Disapproved Withdraw Fwd/ESDIS ERB Fwd/ESDIS

ADDITIONAL SHEET

CCR #: 02-0981 Rev: - Originator:

Telephone: Office:

Title of Change:

1. Install the new 10x180GB drive shelf in the StorageTek 9176 RAID rack.

2. Configure the new drives as two 9+1 RAID 5 LUNs:

2.1 Execute /opt/SM8/client/SMclient

2.2The new disks should be visible in the GUI as "Unconfigured Capacity." Note the total size of this available capacity - it should be close to 169*10 = 1,690GB.

2.3 Right click on the unformatted capacity icon and select "create volume."

In the default host type dialog box, select "Solaris (Sparc)" and click OK.

In the "Create Volume Wizard - Introduction" dialog box, select "Unconfigured Capacity" and click Next.

In the "Create Volume Wizard - Specify Volume Group Parameters" dialog box, select RAID 5 and manual drive selection. Select all 10 drives shown in the box (these should be the new drives) and click apply, then next.

In the "Create Volume Wizard - Specify Volume Parameters" dialog box, set the new volume capacity to be half of the size shown (approx 1690/2 GB - see step 2.2). Select the Customized settings and click next.

In the "Create Volume Wizard - Specify Advanced Volume Parameters" dialog box, select custom I/O characteristics and set the cache read ahead multiplier to 8 and the segment size to 64. Then click finish.

In the "Create Volume Wizard - Creation" dialog box, select "same volume group" and click Yes.

In the "Create Volume Wizard - Specify Volume Parameters" dialog box, leave the new volume capacity unchanged (should be). Select the Customized settings and click next.

In the "Create Volume Wizard - Specify Advanced Volume Parameters" dialog box, set the values as above and click Finish.

In the "Create Volume Wizard - Creation" dialog box, click No to complete the volume creation process.

In the Array Management window, click the "Mappings View" tab.

Right-click the "Undefined Mappings" icon and select "define additional mapping."

In the "Define Additional Mapping" dialog box, select the top level host group. Next, select one of the two unmapped volumes and click Add. The Lun number shown will be assigned to this volume. Next, select the last volume and click Add.

At this point the volumes have been created and assigned to LUN numbers. Since the volumes are in the top level host group, they are now visible to the SAN hosts.

3. Backup any data in /datapool that needs to be available after the filesystem reconfiguration. All data resident in the filesystem will be lost when the new filesystem is built on the existing disks.

4. Unmount the SAN filesystem (n0sas01:/datapool) on all NFS clients.

5. Unshare /datapool from n0sas01.

6. Unmount /datapool on n0sas01.

7. Shutdown all SAN client hosts (n0sas01, n0dps01, n0dig06, n0acg01, n0drg01). These hosts must be rebooted to build devices for the new LUNs (reboot the Suns with the -r option).

8. Run the format command on n0sas01 after the reboot and format the 2 new LUNs with all available space on slice 6.

9. Modify /etc/opt/LSCsamfs/mcf as follows:

** Original MCF file **

```
qfs1          10     ma     qfs1
/dev/dsk/c4t0d0s6      11     mm     qfs1    on     /dev/rdsk/c4t0d0s6
/dev/dsk/c3t0d1s6      12     mr     qfs1    on     /dev/rdsk/c3t0d1s6
/dev/dsk/c4t0d2s6      13     mr     qfs1    on     /dev/rdsk/c4t0d2s6
```

** New MCF file **

```
qfs1          10     ma     qfs1
/dev/dsk/c4t0d0s6      11     mm     qfs1    on     /dev/rdsk/c4t0d0s6
/dev/dsk/c3t0d1s6      12     mr     qfs1    on     /dev/rdsk/c3t0d1s6
/dev/dsk/c4t0d2s6      13     mr     qfs1    on     /dev/rdsk/c4t0d2s6
/dev/dsk/c3t0d3s6      14     mr     qfs1    on     /dev/rdsk/c3t0d3s6
/dev/dsk/c4t0d4s6      15     mr     qfs1    on     /dev/rdsk/c4t0d4s6
```

10. Rebuild the QFS filesystem by executing the following command:

/opt/LSCsamfs/sbin/sammkfs -a 576 qfs1

11. Mount and share /datapool on n0sas01. (Note, for QFS version 3.5.0-41A, the writebehind=1152 option should be added to the mount options. This will avoid a data corruption bug in that version of QFS).

12. Set the ownership of the new disk devices in Sanergy. Execute /opt/Sanergy/config and set the ownership by clicking the "owner" link and entering the hostname in the text field. Leave the tag field blank. Note: due to a Sanergy bug, the fully qualified hostname must be trunkated temporarily during this process.

13. Restore /datapool client mounts*. After mounting, for each client, invoke the Sanergy GUI (/opt/SANergy/config on Solaris and /usr/SANergy/config on IRIX) and re-fuse the /datapool filesystem in Sanergy: Check the checkbox next to /datapool in the mount point list and click the "Set" button.

*Note: a recent trouble ticket from GSFC (SMC000000004948) indicates a problem with filesystem sizes being reported incorrectly for Sanergy clients. The remedy for this is to add the following mount options in the vfstab or fstab files for the /datapool filesystem: acregmin=0,acregmax=0,actimeo=0. This fix is published in technical directive 02-004.

14. Restore any data to /datapool that had been backed up.